



In-system Debugging of PCIe® Devices

Philippe Legros
PCI Express Architect
PLDA



Disclaimer



Presentation Disclaimer: All opinions, judgments, recommendations, etc. that are presented herein are the opinions of the presenter of the material and do not necessarily reflect the opinions of the PCI-SIG®.

Summary



- **Introduction**
- **Common problems in hardware**
- **Debug tools & techniques**
- **Conclusion**

- **Debugging issues in a PCI Express[®] system is often challenging and very time consuming**
 - PCI Express is a complex protocol
 - At least two devices with each a controller and a PHY are involved
 - There is often little or no visibility
 - Problem can span over a long time period

- **But fortunately PCI Express users benefit from a wide ecosystem of:**
 - Debug tools
 - Ways to increase visibility inside a device
 - Methods to troubleshoot issues
 - .. And also experts to assist with debugging

Summary



- Introduction
- **Common problems in hardware**
- Debug tools & techniques
- Conclusion

Common Problems in Hardware



- **Most of the problems seen in hardware are:**
 - Link is not training
 - Link is training but speed / width is not as expected
 - Missing/unexpected packets
 - Unexpected entry in Recovery

Link Training Issues



- **Common problems are:**
 - No link is ever established
 - Reset issue
 - Receiver detection issue
 - Link quality issue
 - Link width is not as expected
 - Lane(s) reliability issue

- **There are two main scenarios:**
 - Device is not ready quickly enough after reset
 - Host system enumerates before device is ready and does not even attempt to train link at highest possible speed.
 - Speed change to 8.0 GT/s or 16 GT/s occurs but there is a timeout in equalization
 - Link automatically backs down to a lower speed (selected speed may depend on platform)

- **This is application specific so there are many possible symptoms:**
 - Missing TLPs
 - Duplicate TLPs
 - Unexpected traffic

- **There are *many* reasons that can cause LTSSM to go to Recovery:**

- Link quality issue
 - Encoding error
 - Framing error
- Flow-control timeout
- Device / link partner malfunction
- Low-power entry/exit issue
- ...



Summary



- Introduction
- Common problems in hardware
- **Debug tools & techniques**
- Conclusion

Debug Tools and Techniques



- **Advanced error reporting**
- **Protocol analyzer**
- **Test interface**
- **Embedded analyzer**
- **Link reducers / Passive interposer card**

Advanced Error Reporting



- **AER is an optional feature of PCIe® protocol that provides detailed information about errors that occur in a device:**
 - Separate status/mask/severity bit for each type of error
 - Offending TLP header is logged when applicable
 - Supports internal (i.e. application specific) errors

Advanced Error Reporting



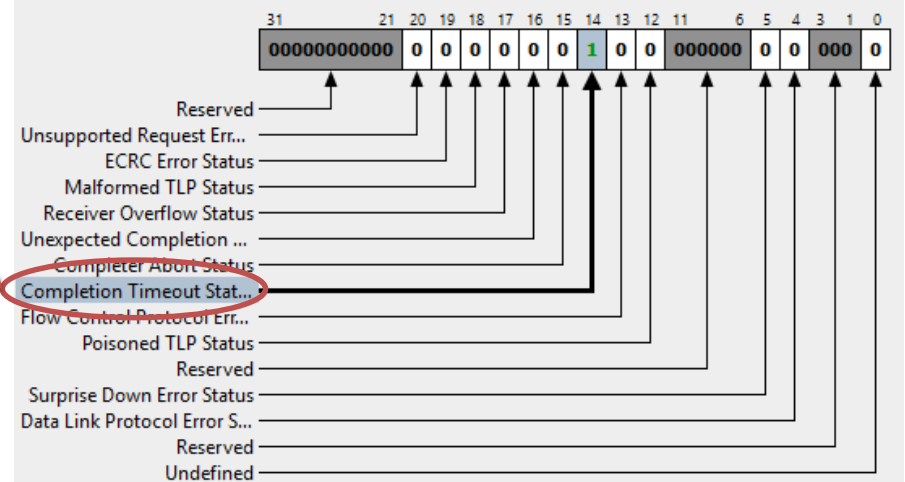
- **Example:** “My device performs R/W DMA to system memory, but it stops transferring data after running for several minutes...”

Advanced Error Reporting

- **AER registers report that a Completion Timeout occurred:**

Advanced Error Reporting Extended Capability Header	100h
14010001	
Uncorrectable Error Status Register	104h
00004000	
Uncorrectable Error Mask Register	108h
00000000	
Uncorrectable Error Severity Register	10Ch
00462031	
Correctable Error Status Register	110h
00000000	
Correctable Error Mask Register	114h
00002000	
Advanced Error Capabilities and Control Register	118h
00000000	
Header Log Register (1st DW)	11Ch
00000000	
Header Log Register (2nd DW)	120h
00000000	
Header Log Register (3rd DW)	124h
00000000	
Header Log Register (4th DW)	128h
00000000	

Uncorrectable Error Status Register
Advanced Error Reporting Capability



Advanced Error Reporting



- **So DMAs stopped because some completions are missing!**
- **Probable causes:**
 - Malformed read request
 - Design error in transmit or receive logic
 - ...

- **This is the best known and most versatile tool to debug PCIe[®] issues**
 - Easy view of traffic
 - Powerful trigger conditions
 - Can filter out specific packets
 - Can record long sequences

Protocol Analyzer

○ Example: Unexpected link retraining

Packet		8.0	TS1	Link	Lane	N_FTS	Training Control	Data Rate	Eq Control	Pre-Cursor	Cursor	Post-Cursor	TS1 Symbols	Time Delta	Time Stamp
Packet 14793	R→	x1	TS1	0	0	255	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 1 0	0	20	4	4A...	15.500 ns	0005 . 095 608 962 s
Packet 14794	R←	x1	TS1	0	0	32	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A...	0.750 ns	0005 . 095 608 978 s
Packet 14795	R→	x1	TS1	0	0	255	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 1 0	0	20	4	4A...	15.500 ns	0005 . 095 608 978 s
Packet 14796	R←	x1	TS1	0	0	32	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A...	0.750 ns	0005 . 095 608 994 s
Packet 14797	R→	x1	TS1	0	0	255	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 1 0	0	20	4	4A...	15.500 ns	0005 . 095 608 994 s
Packet 14798	R←	x1	TS1	0	0	32	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A...	0.750 ns	0005 . 095 609 010 s
Packet 14799	R→	x1	TS2	0	0	255	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 1 0	45 ...				15.500 ns	0005 . 095 609 010 s
Packet 14800	R←	x1	TS1	0	0	32	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A...	0.750 ns	0005 . 095 609 026 s
Packet 14801	R→	x1	TS2	0	0	255	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 1 0	45 ...				15.500 ns	0005 . 095 609 026 s
Packet 14802	R←	x1	TS1	0	0	32	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A...	0.750 ns	0005 . 095 609 042 s
Packet 14803	R→	x1	TS2	0	0	255	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 1 0	45 ...				15.500 ns	0005 . 095 609 044 s
Packet 14804	R←	x1	TS1	0	0	32	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A...	0.750 ns	0005 . 095 609 058 s
Packet 14805	R→	x1	TS2	0	0	255	0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 1 0	45 ...				15.500 ns	0005 . 095 609 060 s

So, how to proceed ?

Protocol Analyzer

- First step is always to check which device initiates retraining

Packet	R+	5.0	DLLP	ACK	AckNak_Seq_Num	CRC 16	Idle	Time Stamp										
14752	R+	x1			11	0x5893	224.382 us	0005 . 095 108 254 s										
Packet	R+	5.0	TLP	Cfg	CfgWr0	Length	RequesterID	Tag	DeviceID	Register	1st BE	Device ID	Vendor ID	LCRC	Time Delta	Time Stamp		
14753	R+	x1	10		010:00100	1	000:00:0	5	001:00:0	0x000	1111	0xFFFF	0xFFFF	0x07EC58B2	263.500 ns	0005 .		
Packet	R+	5.0	DLLP	ACK	AckNak_Seq_Num	CRC 16	Idle	Time Stamp										
14754	R+	x1			10	0xF985	114.000 ns	0005 . 095 332 908 s										
Packet	R+	5.0	TLP	Cpl	CplID	Length	RequesterID	Tag	CompleterID	Status	BCM	Byte Cnt	Lwr Addr	LCRC	Time Delta	Time Stamp		
14755	R+	x1	12		000:01010	0	000:00:0	5	001:00:0	SC	0	4	0x00	0x852946CE	536.750 ns	0005 . 0		
Packet	R+	5.0	DLLP	ACK	AckNak_Seq_Num	CRC 16	Idle	Time Stamp										
14756	R+	x1			12	0x3FD1	262.078 us	0005 . 095 333 566 s										
Packet	R+	5.0	TLP	Cfg	CfgWr0	Length	RequesterID	Tag	DeviceID	Register	1st BE	LCRC	Time Delta	Time Stamp				
14757	R+	x1	11		000:00100	1	000:00:0	7	001:00:0	0x000	1111	0x5429E605	271.500 ns	0005 . 095 595 654 s				
Packet	R+	5.0	DLLP	ACK	AckNak_Seq_Num	CRC 16	Idle	Time Stamp										
14758	R+	x1			11	0x5893	109.750 ns	0005 . 095 595 924 s										
Packet	R+	5.0	TLP	Cpl	CplID	Length	RequesterID	Tag	CompleterID	Status	BCM	Byte Cnt	Lwr Addr	Device ID	Vendor ID	LCRC		
14759	R+	x1	13		010:01010	1	000:00:0	7	001:00:0	SC	0	4	0x00	0x1100	0x1556	0xC0C58937		
Packet	R+	5.0	TLP	Cpl	CplID	Length	RequesterID	Tag	CompleterID	Status	BCM	Byte Cnt	Lwr Addr	Device ID	Vendor ID	LCRC		
14760	R+	x1	13		010:01010	1	000:00:0	7	001:00:0	SC	0	4	0x00	0x1100	0x1556	0xC0C58937		
Packet	R+	5.0	TLP	Cpl	CplID	Length	RequesterID	Tag	CompleterID	Status	BCM	Byte Cnt	Lwr Addr	Device ID	Vendor ID	LCRC		
14761	R+	x1	13		010:01010	1	000:00:0	7	001:00:0	SC	0	4	0x00	0x1100	0x1556	0xC0C58937		
Packet	R+	5.0	TLP	Cpl	CplID	Length	RequesterID	Tag	CompleterID	Status	BCM	Byte Cnt	Lwr Addr	Device ID	Vendor ID	LCRC		
14762	R+	x1	13		010:01010	1	000:00:0	7	001:00:0	SC	0	4	0x00	0x1100	0x1556	0xC0C58937		
Packet	R+	5.0	EIEOS	EIEOS Symbols	Idle	Time Stamp												
14763	R+	x1		00 FF 00 ...	0.250 ns	0005 . 095 608 558 s												
Packet	R+	5.0	TS1	Link Lane N_FTS Training Control	Data Rate	Eq Control	Pre-Cursor	Cursor	Post-Cursor	TS1 Symbols	Idle	Time Stamp						
14764	R+	x1		0 0 32 0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A ...	0.250 ns	0005 . 095 608 576 s						
Packet	R+	5.0	TS1	Link Lane N_FTS Training Control	Data Rate	Eq Control	Pre-Cursor	Cursor	Post-Cursor	TS1 Symbols	Idle	Time Stamp						
14765	R+	x1		0 0 32 0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A ...	0.250 ns	0005 . 095 608 592 s						
Packet	R+	5.0	TS1	Link Lane N_FTS Training Control	Data Rate	Eq Control	Pre-Cursor	Cursor	Post-Cursor	TS1 Symbols	Idle	Time Stamp						
14766	R+	x1		0 0 32 0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A ...	0.250 ns	0005 . 095 608 608 s						
Packet	R+	5.0	TS1	Link Lane N_FTS Training Control	Data Rate	Eq Control	Pre-Cursor	Cursor	Post-Cursor	TS1 Symbols	Idle	Time Stamp						
14767	R+	x1		0 0 32 0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A ...	0.250 ns	0005 . 095 608 624 s						
Packet	R+	5.0	TS1	Link Lane N_FTS Training Control	Data Rate	Eq Control	Pre-Cursor	Cursor	Post-Cursor	TS1 Symbols	Idle	Time Stamp						
14768	R+	x1		0 0 32 0 0 0 0	2.5 GT/s, 5 GT/s, 8 GT/s	0 0 0 0	0	30	10	4A ...	0.250 ns	0005 . 095 608 640 s						

- **Then check if this device:**
 - Receives incorrect packets
 - Sends error messages (can be after retraining)
 - Is not receiving expected ACK, flow control DLLP, completion,...

-
- The screenshot displays a network traffic analysis tool interface. The main window shows a list of captured packets. The left pane lists packets with their sequence numbers, times, and protocols. The right pane shows the details of the selected packet, including its structure and data. The bottom pane shows the raw data of the selected packet. Red arrows highlight specific packets and their details.
- | Packet | Time | Protocol | Source | Destination | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC | Device ID | Vendor ID | LCRC | Time Stamp |
|--------|------|----------|-------------|-------------|------------------|------------|------------|---------------|--------|------------|----------|--|-----------|-----------|------|------------|
| 14694 | 0.0 | TSP2 | Link | Line N_FTS | Training Control | Data Rate | Eq Control | TSP2 Symbols | Idle | Time Stamp | | | | | | |
| 14695 | 0.0 | TSP2 | Link | Line N_FTS | Training Control | Data Rate | Eq Control | TSP2 Symbols | Idle | Time Stamp | | | | | | |
| 14696 | 0.0 | SOS | SOS Symbols | Time Stamp | | | | | | | | | | | | |
| 14697 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14698 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14699 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14700 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14701 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14702 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14703 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14704 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14705 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14706 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14707 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14708 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14709 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14710 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14711 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14712 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14713 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14714 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14715 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14716 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14717 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14718 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14719 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14720 | 0.0 | TCP | Meg | Meg Routing | Length | Request ID | Tag | Completion ID | Status | BCM | Byte Cnt | LCRC <td>Device ID</td> <td>Vendor ID</td> <td>LCRC</td> <td>Time Stamp</td> | Device ID | Vendor ID | LCRC | Time Stamp |
| 14721 | | | | | | | | | | | | | | | | |

22

- **Protocol analyzers have limitations:**

- They can show what the problem is, but not what the reason is !
- They have limited ability to follow LTSSM and can show incorrect information.
- They do not necessarily “see” all traffic :



- Example : at 8.0 GT/s or higher some traffic can be missed (Equalization phase 0/1 for example) if analyzer does not lock on data quickly enough !

- **Having access to internal status signals can give precious hints:**
 - LTSSM
 - Error flags
 - Status information
 - Event counters
 - ...

These become even more precious if link cannot be established or is lost

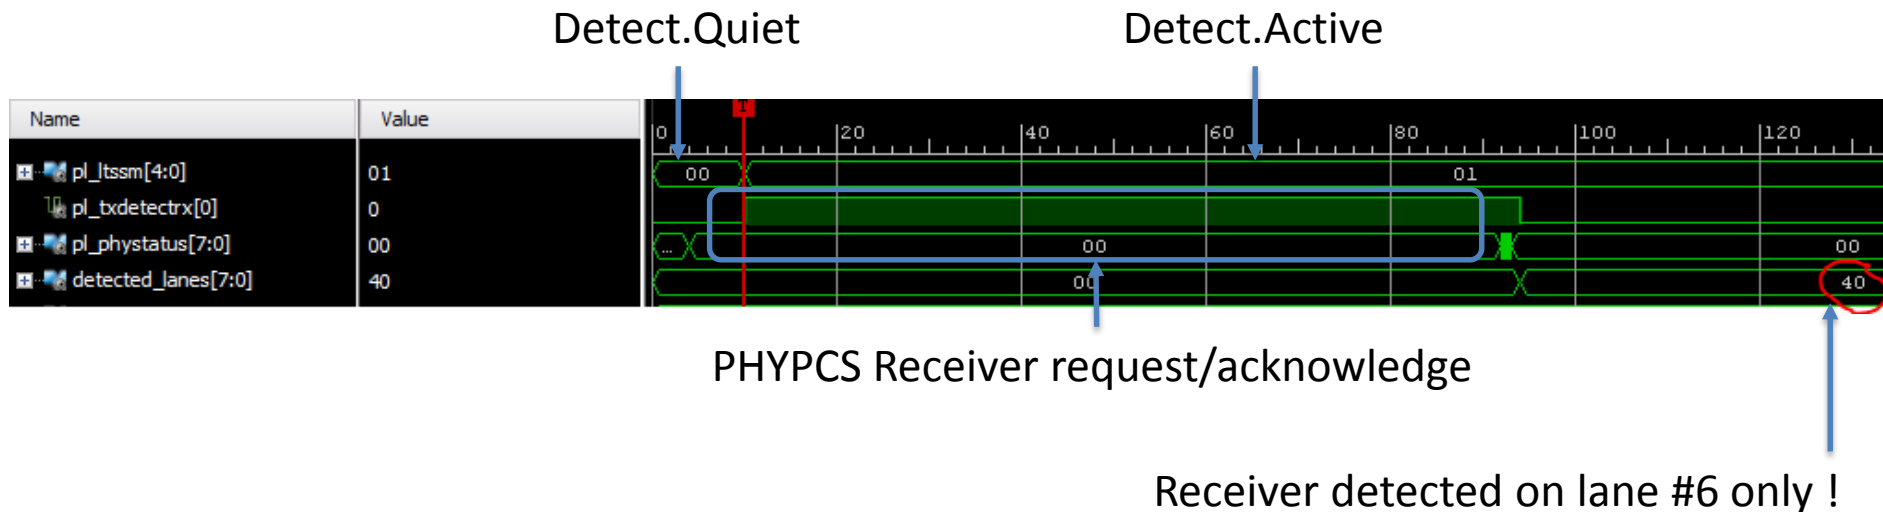
- **These tools can provide unparalleled insight on device:**
 - Internal states such as LTSSM
 - Error & event flags
 - Activity/traffic indicators
 - PHYPCS interface

- **Easy to use and widely available**

Embedded Analyzer

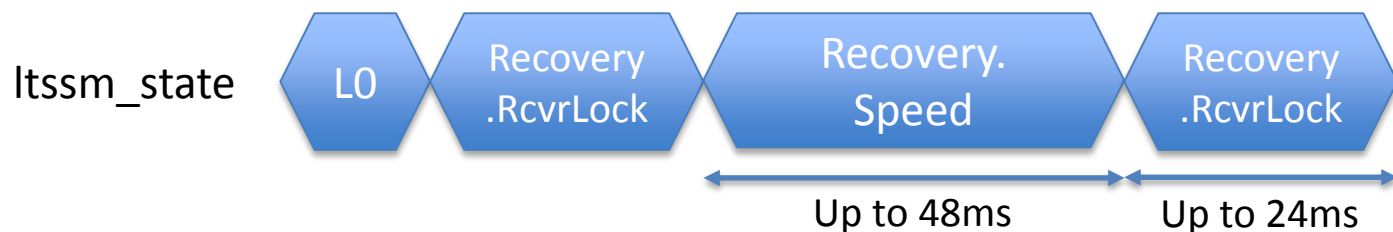


- **Example: Checking that receiver detection is behaving properly**



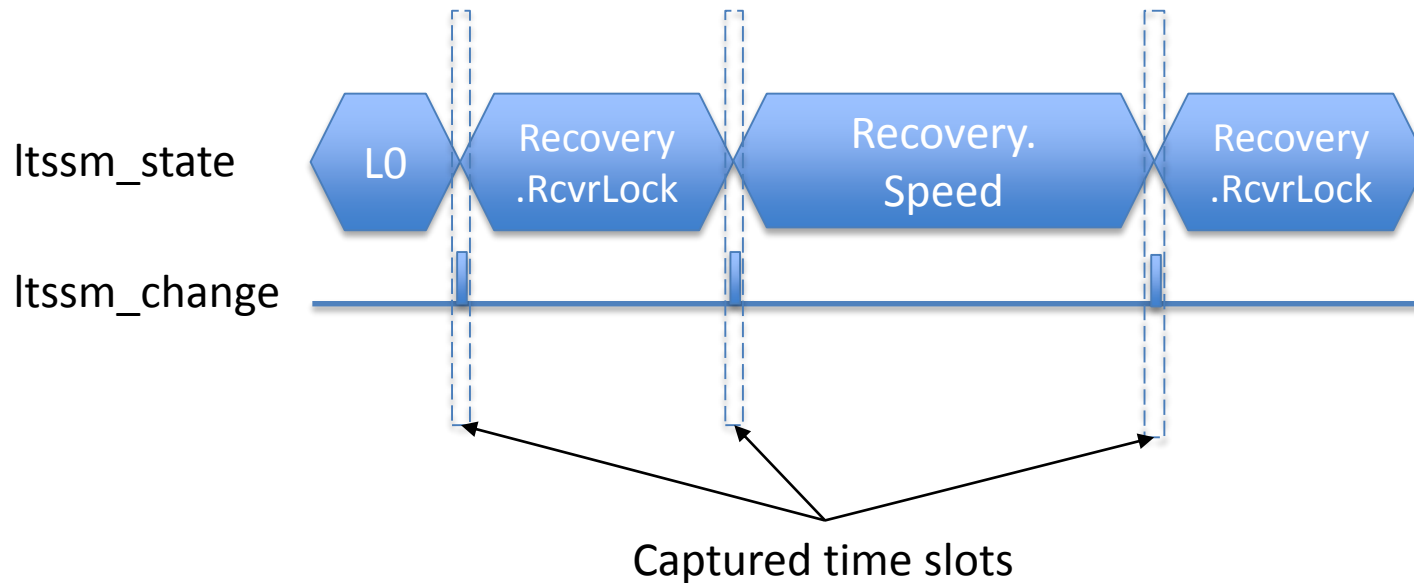
- **It is often difficult to have an overview of link training process:**
 - Transition between states can take several milliseconds
 - Timeout values in specification are 2...48 ms

Example: Speed change



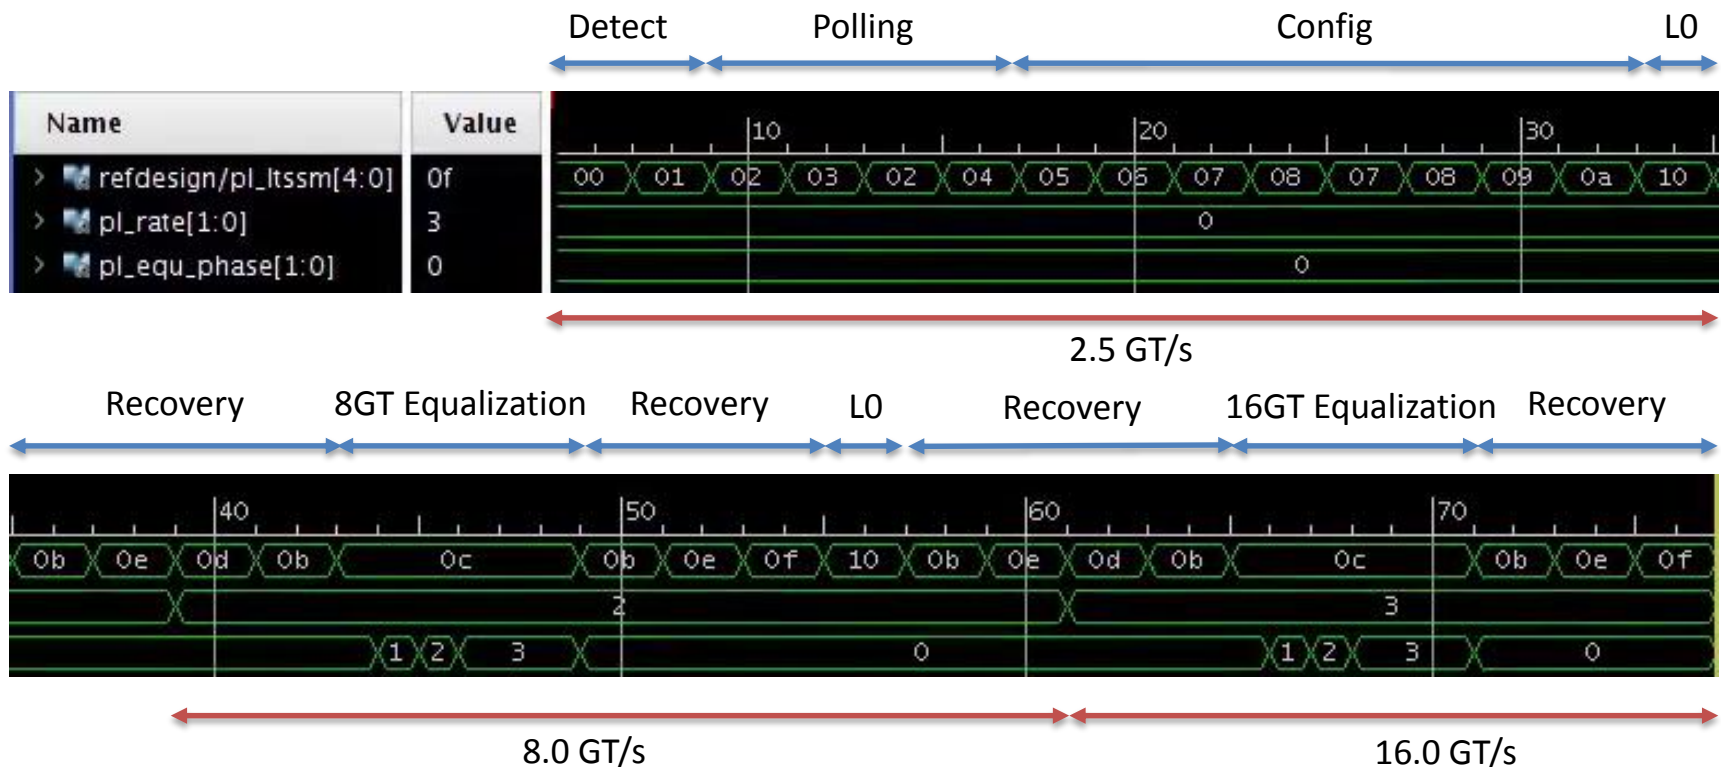
.. and embedded analyzer capture window is very limited

- This can be solved by using a “storage qualifier” signal to tell embedded analyzer when to store data

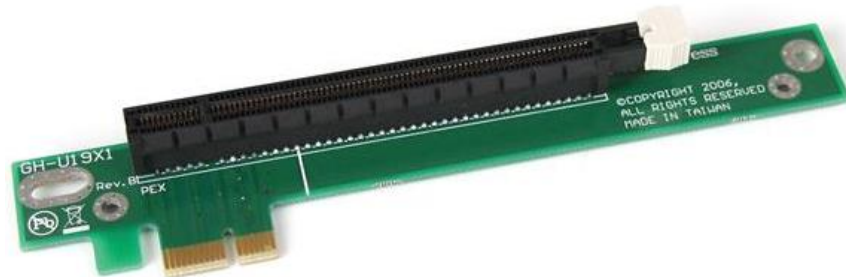


Embedded Analyzer

- A complete link training sequence can then be captured with minimal sample buffer size:



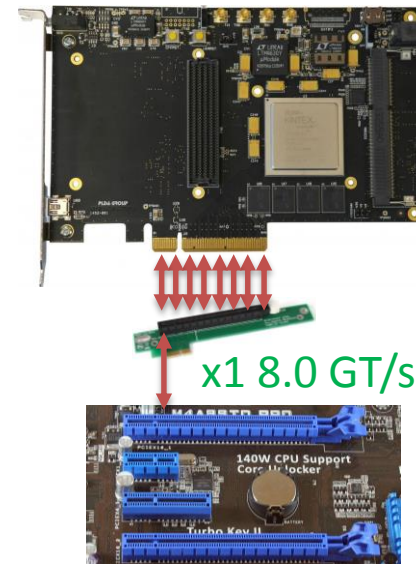
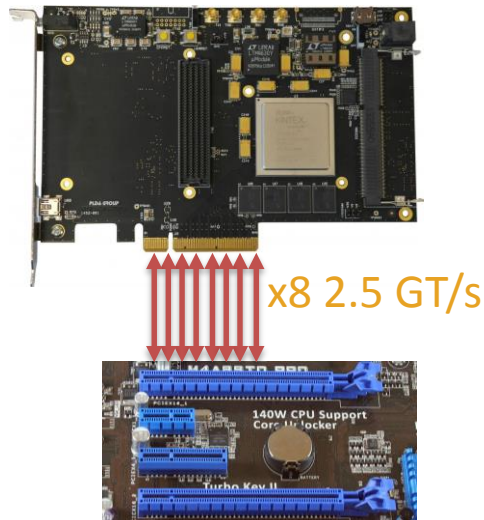
- **Sometimes the most inexpensive accessory can be very handy and help troubleshoot nasty issues..**



Should be used with caution: link reducers can sometimes deteriorate link quality and thus change device/system behavior

Link Reducer

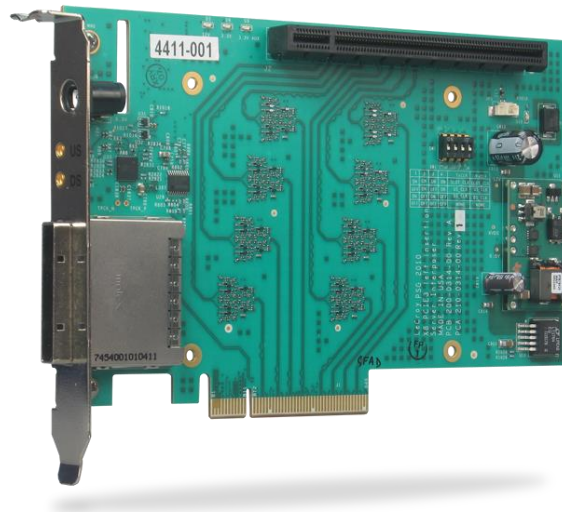
- “I have a x8 8GT/s device, but it only comes up in x8 2.5GT/s in our platform..”
 - One quick experiment is to try with a link reducer:



- Link comes up at 8.0 GT/s in x1, showing that there is probably a link reliability issue on one or more lanes

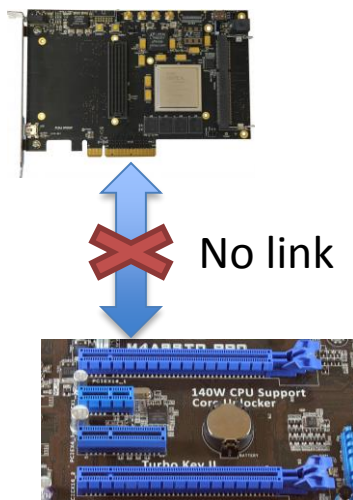
Passive Interposer Card

- **Normally comes with a PCIe protocol analyzer, however it can also be a valuable stand-alone tool**
 - Can help troubleshoot some electrical issues



Passive Interposer Card

- “My device works in some systems, but in some others there is no link at all.”
 - A simple test is to plug it on a passive interposer card:



- Link comes up with an interposer, showing that there is an electrical problem (impedance,...)

Summary



- Introduction
- Common problems in hardware
- Debug tools & techniques
- **Conclusion**

- **There is not a single tool that can help troubleshoot all in-system issues**
 - It is important to understand which tool is the best suited for each type of problem
 - Having multiple tools / sources of information never hurts

- **Planning debugging during product development is absolutely critical!**
 - Enabling AER
 - Selecting key signals and adding them to a test port or embedded logic analyzer
 - Having ability to enable/disable features..

Each of these can be a life-saver



**Thank you for attending the
PCI-SIG Developers Conference 2017.**

**For more information please go to
www.pcisig.com**

